

The Antitrust Bulletin

The Journal of American and Foreign
Antitrust and Trade Regulation

**Microsoft plays hardball:
the use of exclusionary pricing
and technical incompatibility
to maintain monopoly power
in markets for operating
system software**

BY KENNETH C. BASEMAN,* FREDERICK R.
WARREN-BOULTON* and GLENN A. WOROCH**

Reprinted from

VOL. XL, NUMBER 2—SUMMER 1995

Microsoft plays hardball: the use of exclusionary pricing and technical incompatibility to maintain monopoly power in markets for operating system software

BY KENNETH C. BASEMAN,* FREDERICK R.
WARREN-BOULTON* and GLENN A. WOROCH**

I. Introduction and summary

This article examines Microsoft's licensing practices for its MS-DOS and Microsoft Windows operating system software. Our

* Principals, MICRA: Microeconomic Consulting and Research Associates, Inc., Washington, DC.

** University of California, Berkeley.

AUTHORS' NOTE: We would like to express our appreciation for helpful comments and other assistance to Sturge Sobin, Linnet Harlan, Paul Dennis and the participants at the Columbia Business School's Institute for Tele-Information's Seminar on Sustaining Competition in Network Industries Through Regulating and Pricing Access, especially Janusz Ordover and Bobby Willig.

main focus is on Microsoft's use of CPU (central processing unit, or per-processor) licenses, under which an original equipment manufacturer (OEM) of personal computers pays a royalty for each machine it ships instead of for each unit of MS-DOS installed. We also examine Microsoft's practice of requiring in these licenses a minimum number of personal computers (PCs) on which MS-DOS can be installed, Microsoft's tying of Microsoft Windows and technical support information to the sale of MS-DOS, and Microsoft's attempts to induce technical incompatibility between MS-DOS and its main competitor, DR-DOS. Finally, we turn to the proposed consent decree between Microsoft and the Department of Justice.

We begin in section II with a brief description of the market in the early 1990s for personal computer operating systems, and a history of Microsoft's licensing practices and technical design tactics. In addition to CPU licensing, the minimum requirements contract, and the Microsoft Windows tie to MS-DOS, we discuss "cliff-pricing" quantity discounts for MS-DOS and attempted or threatened design incompatibilities between MS-DOS and DR-DOS.

We then turn to two potential efficiency rationales for the CPU license. The first, presented in section III, is based on the similarity between a CPU license and a two-part tariff, which can result in "efficient" first-degree price discrimination. Upon closer examination, however, we find that the CPU license could operate as a two-part tariff only toward customers with time horizons of less than 1 year. Moreover, even if a CPU license were a two-part tariff, in this specific factual context a CPU license would be neither welfare enhancing nor—absent an exclusionary effect—would it even be profit-maximizing.

Section IV examines a second potential efficiency rationale for the CPU: reducing OEM fraud and/or software piracy. By reducing the number of "naked" machines shipped by OEMs, a CPU license could deter OEMs from engaging in fraud, or it could deter OEMs, retailers and/or customers from piracy. An examination of the historical record leads us to conclude, however, that the prevention of piracy and fraud is not a plausible

explanation for why CPU licensing was introduced. Even more telling, however, is that the CPU license is no more effective at deterring piracy or fraud than are other available but unused nonexclusionary alternatives such as a "credited-CPU" license.

Section V turns to potential anticompetitive rationales for Microsoft's practices in the DOS market. We begin by observing that markets for many high technology products are characterized by a competitive process where a new product appears with a significantly superior technology or design and sweeps the field. By rapidly displacing the old product and its old technology, it achieves a very high market share in a very short time and earns very large profits. This situation persists only until the dominant firm's product is itself displaced by a new superior product. This cycle of a new product with a new technology displacing an existing product with an old technology is a process of "creative destruction" in the race to be best. Firms achieve a dominant position, but hold that position only transitorily because, without artificial barriers to entry, today's dominant or monopoly firm and product can readily be displaced by a new product developed by a competitor or a new entrant.

When the monopolist's position is protected by strategically erected barriers to entry, however, this displacement process can come to a halt. We examine the possibility that Microsoft has used a variety of exclusionary practices, notably nonlinear pricing and technical incompatibility, not to achieve its initial position but rather to retain that position against new competition. We conclude that, under the conditions present in the operating systems market, such practices can be, and in this instance have been, effective in limiting the growth and threatening the existence of entrants and rivals with very small market shares. We also conclude that Microsoft's anticompetitive behavior has reduced social welfare.

In section VI we turn to an examination of the Department of Justice's recent settlement with Microsoft. While we do not quarrel with the Department's focus on the horizontal aspects of the case, we are concerned that the remedies prescribed in the consent decree are likely to be inadequate. Specifically, the consent decree

fails to prevent Microsoft from employing nonlinear pricing or quantity discounts to achieve the same exclusionary consequences as the offending practices. We offer several remedies—including a ban on sales or discounting of naked machines, “credited CPU licenses” and allowing arbitrage—that lack the exclusionary aspects of CPU licenses. Furthermore, these alternatives preserve any antipiracy and antifraud properties that CPU licenses may possess.

II. Background

A. The market for personal computer operating systems

1. PERSONAL COMPUTER PLATFORMS Our focus is on the market for packaged software that operates personal computers, and to a lesser extent, the software applications that run using those operating systems. To better understand the market for these products, we must delve into the economics and technology of the personal computer.

Several key features distinguish this product. PCs can be decomposed into hardware and software components. Some of these components are essential: every computer system requires a microelectronic chip called the central processing unit (CPU), plus operating system (OS) software. The OS directs the stream of instructions requested by the applications software, while the CPU performs the numerical computations. Importantly, the CPU and the OS are almost always combined in fixed proportions: one of each is needed per system.

Once an OS is installed, a user can run many kinds of applications software.¹ The most popular packages do word processing, spreadsheet analysis, and database management. Increasingly popular is the use of graphical user interface (GUI) that simplifies the management of the various applications. Both applications and GUIs are optional components of a personal computer system.

¹ A PC also requires a layer of software that stands between the CPU and the OS. Called the BIOS, or basic instruction operating system, this code is burned into the machine’s ROM (read only memory) chip.

Personal computers are available in several "platforms" that differ in their hardware specifications. The predominant PC "platform" is the so-called IBM-compatible PC, which has evolved from the hardware and software specifications of the machine introduced by IBM in 1981.

2. INDUSTRY STRUCTURE The supply of many components is highly concentrated. First of all, an overwhelming proportion of IBM-compatible PCs in use today are equipped with CPUs manufactured by the Intel Corporation. Secondly, the majority of existing PCs run on one version or another of the operating system sold by Microsoft Corporation. Sales of applications software and peripheral hardware components are far less concentrated.²

Hundreds of OEMs assemble hardware components in various configurations called "models," distribute the machines through retail stores or mail order, and provide technical and repair service. In addition to a few large OEMs such as Compaq, Dell and AST Research in the U.S., and NEC, Toshiba and Hitachi in Asia, there is a host of small resellers. We can safely assume this segment of the market to be competitive.

The bulk of new PCs shipped in the U.S. (see the table) arrive loaded with some operating system, usually Microsoft's MS-DOS, and often with their Microsoft Windows interface as well. IBM ships its PCs with one of its own operating systems: PC-DOS or OS/2.³ The only independent DOS (i.e., compatible with, but not a clone or derivative of MS-DOS) was Digital Research Incorporated (DRI's) DR-DOS which, with Novell's acquisition of DRI in 1991, is now referred to as Novell DOS and IBM's PC-DOS. Users could purchase OSs at retail stores (e.g., Egghead Software) or direct from the software publisher.

3. DEMAND FOR PCs Personal computers are purchased by businesses, individuals, schools, government, and other organizations. Users select a hardware platform and a specific hardware-software

² As to the ROM-BIOS, there are many providers including Phoenix, AMI, Quadtel and Award in addition to IBM's original version.

³ OS/2 combines OS software and a GUI in one program.

bundle based on the overall system price. Accordingly, demand for an operating system as well as for the microprocessor is derived from demand for the entire package. In 1992, it was estimated that the worldwide installed base of personal computers of all platforms totaled over 138 million.⁴ Of those, 72% were IBM-compatible. Less than a quarter of those machines were equipped with Microsoft Windows.

Potential purchasers of an operating system can be divided into two groups. First, there are existing PC owners seeking to upgrade their OS or switch to a new one. Then there are individuals who need an OS for a new machine, whether it is their first PC, a replacement, or an additional one. Either way, current OS users will bear some costs when switching from one OS to another. Certainly this would be true if the new OS demands a more powerful machine (as with the move from DOS to either Microsoft Windows⁵ or OS/2). At a minimum, the user must learn some new command or menu structure and may have to replace outdated or incompatible applications.

4. SUPPLY CONDITIONS Operating system software is very costly to develop and market. For instance, it has been estimated that IBM has spent over \$2 billion developing OS/2. In comparison, operating system software is relatively cheap to produce and maintain. As a result, fixed costs are enormous and marginal costs are negligible. These fixed costs are also largely sunk. The code itself is rarely of little value in other uses.⁶ Development teams accumulate expertise and reputation, only a portion of which can be redeployed into other projects.

⁴ BERNSTEIN RESEARCH, FTC INVESTIGATION OF MICROSOFT (Sanford C. Bernstein & Co.: New York, January 1993).

⁵ Strictly speaking, Windows is not an OS, in that it must run in conjunction with DOS.

⁶ This may change as operating systems adopt the object-oriented approach in which the program is composed of "objects" that can be reused in other programs in a modular fashion. NeXT Computer's operating system, NextStep, has pursued this strategy.

Entry into the operating systems software market is not easy. There are sunk costs of development and marketing, as mentioned above. Besides the irreversible investment in computer code, incumbents acquire sunk, or partially sunk, assets such as customer lists and brand name recognition. Furthermore, any new OS must be compatible with all the applications that were written to that "standard." User switching costs also limit the ability of new entrants to gain a toehold. Of course, these costs erect barriers only when the incumbent firm has a first-mover advantage. However, sunk costs ordinarily imply a first-mover advantage, at least for the current vintage of technology.⁷ In addition, as we will see, IBM bestowed a somewhat unique first-mover advantage on Microsoft when it selected MS-DOS to be the operating system for its PC.

5. HISTORY OF PC OPERATING SYSTEMS Dating back to 1976, Digital Research Incorporated sold a popular operating system, called CP/M, for use on machines based on Intel's 8-bit 8080 chip. In 1980, in what may become the deal of the century, Microsoft paid a mere \$100,000 for the rights to a CP/M derivative or clone software package called "Disk Operating System," which, with minor modifications, became the initial MS-DOS. In 1981, when IBM launched its entry into the personal computer market, it selected Intel's new 16-bit 8088 chip as the CPU. It also chose to endorse Microsoft's MS-DOS as the operating system.

IBM's partnership with Microsoft later fell apart. But in the meantime, neither IBM nor DRI stopped developing their own operating systems.⁸ Under the terms of the dissolution, IBM continued to develop MS-DOS, and eventually its own variant,

⁷ If there are no cost complementarities across vintages of technology, then the requirement to sink substantial investment in software development will not convey an advantage to the successful first generation firms in the competition to develop subsequent generations of technology.

⁸ There were two other significant MS-DOS derivatives. For a while, Compaq Computer had shipped its machines with its own Compaq DOS, and NEC developed NEC-DOS, a proprietary operating system that, until recently, dominated the Japanese market.

PC-DOS, which it loaded on PCs bearing the IBM nameplate. In exchange, IBM agreed to pay Microsoft a royalty for a predetermined number of units.

Having been passed over by IBM, DRI went on to modify CP/M for the Intel 8086 chip, leading to its CP/M-86. Later it developed DOS PLUS and then DR-DOS. In April 1990, DRI introduced DR-DOS 5.0 to critical acclaim. Instantly, it began to make inroads into MS-DOS 4.0's market share. By year-end 1990, DR-DOS's share had increased to 10% of new OS shipments, leaving MS-DOS with 70% and IBM with 18%.⁹

Within a month of DR-DOS 5.0's inauguration, Microsoft reported development of MS-DOS 5.0. Curiously, it boasted nearly all of the innovative features of the DRI product. Yet MS-DOS 5.0 was not commercially available until July 1991, more than a year after DR-DOS 5.0's release. Anticipation of the new Microsoft product, prolonged by continuous Microsoft statements indicating imminent availability, however, reined in growth of DR-DOS 5.0 sales.¹⁰

The emergence of the graphical interface played an important role in the events that followed. After repairing bugs in Microsoft Windows 3.0, Microsoft shipped Microsoft Windows 3.1 in April 1991. In that year, 18.5% of new PC shipments included Microsoft Windows along with MS-DOS. By 1992, that fraction jumped to 59.7%. Over that period, sales of MS-DOS (both with and without Microsoft Windows 3.1) rose 28.9% while sales of PC-DOS and DR-DOS fell 15.4%. (See table.) By 1993, the market shares for operating systems on x86 PCs were 79% for MS-DOS, 13% for PC-DOS, 4% for OS/2, 3% for DR-DOS and 1% for UNIX.¹¹

⁹ See BERNSTEIN RESEARCH, *supra* note 4, exhibit 2.

¹⁰ Sherer, *Microsoft Outlines DOS 5.0 to Ward Off DR-DOS*, PC WEEK, October 22, 1990, at 10.

¹¹ See note 28 *infra*.

Table

New Shipments of Personal Computer Operating Systems

Company	Operating system	1990	1991	1992
Microsoft	MS-DOS	11,648	13,178	18,525
	w/Windows	490	2,440	11,056
	w/o Windows	11,158	10,738	7,469
IBM	PC-DOS	3,031	3,003	2,315
DRI/Novell	DR-DOS	1,737	1,819	1,617
DOS Subtotal		16,603	18,288	22,847
Apple	Macintosh	1,411	2,204	2,570
UNIX	UNIX	357	582	797
IBM	OS/2	0	0	409
Other	NEC, etc.	5,079	4,628	4,458
TOTALS		23,450	25,702	31,080

SOURCE: Bernstein Research, International Data Corporation.

B. Microsoft's practices

1. THE CPU LICENSE When first available, MS-DOS was sold to OEMs for a flat fee. Microsoft offered an unlimited number of copies for \$95,000, and for a limited time, reduced the price by half.¹² Around 1983, Microsoft began to gear its license fees to the level of OEM sales. Then and now, each OEM contract was individually negotiated; an external price list never existed.

Over time, Microsoft phased in a new type of royalty contract. By 1992, the "CPU license" became the dominant sales arrangement, with 60% of Microsoft's operating system sales made under CPU licenses.¹³ Under its terms, affiliated OEMs were required to

¹² MANES & ANDREWS, GATES: HOW MICROSOFT'S MOGUL REINVENTED AN INDUSTRY—AND MADE HIMSELF THE RICHEST MAN IN AMERICA (1993); cites are to edited and condensed version in *Microsoft Monopoly*, UPSIDE, March 1993, at 10–18.

¹³ The percentage of Microsoft's operating system sales made under CPU agreements rose from 20% in FY 1989 to 22% in FY 1990, 27% in FY 1991 and 50% in FY 1992. By FY 1993, 60% of MS-DOS sales to OEMs and 43% of Windows sales to OEMs were covered under CPU agreements. See note 28, *infra*.

pay a royalty for every CPU they shipped. Since each machine had a single CPU, the OEM paid for a copy regardless of whether the machine was preloaded with MS-DOS. Microsoft would sell DOS licenses to OEMs who refused the CPU license, but only at significantly higher prices.

Under the CPU license, an OEM usually had to also commit to a minimum "requirement" (X) that approximates its annual shipments. The one-time charge for this requirement is computed using a negotiated per-unit price (f) multiplied by X .¹⁴ If an OEM shipped a machine with a competing operating system, say PC-DOS or DR-DOS, *it would receive no reduction in its payment to Microsoft*. Consequently, an OEM who accepts a CPU license faces a zero marginal price for units of MS-DOS up to the minimum requirement. In the event an OEM exceeded its projected volume during the contract period, the per-unit fee (f) used to calculate the lump sum payment for the first X units would apply to each unit above X . Thus, once the contract is in place, the marginal price is 0 up to X units and f for additional units.

Regardless of whether an OEM ends up shipping more or less than X PCs during the contract, the terms of the CPU license commit the OEM to pay for one unit of MS-DOS for each PC it ships. As a result, customers view themselves as paying double if they use other OSs. If an OS competitor offers to sell at a per-unit price m , the OEM will only buy if the second OS has a quality advantage over MS-DOS valued at m or more.¹⁵

¹⁴ It is paid to Microsoft over the course of the year with an initial payment at the beginning of the year.

¹⁵ For a competing OS supplier to make a sale to an OEM who has signed the CPU license, the quality differential must be worth at least m regardless of whether the OEM is shipping more or less than X PCs. When the OEM ships less than X units, her marginal cost of using MS-DOS on the next PC is zero, compared with a marginal cost of m if she chooses another OS. After X PCs have been shipped, the marginal OS cost is f if the OEM uses MS-DOS, and it is $f + m$ if the OEM uses the other OS. In each case, the additional marginal cost of using the alternative DOS is m . If the machine is shipped "naked," then m is zero.

In 1992, the average license fee per copy of MS-DOS to a hardware OEM under these CPU licenses has been estimated at \$15, far below the average retail price of an upgrade of \$49.¹⁶ All together in that year, Microsoft grossed \$399 million on world-wide sales of 18,525,000 units of MS-DOS to OEMs and as upgrades.¹⁷ From every indication, the implicit per-unit charges and requirement levels vary across the contracts signed by different OEMs.

Typically, these agreements ran for a period of 2 years. It was quite likely an OEM will finish any contract with unused licenses—if only because Microsoft, in an attempt to lock-in customers, would offer a lower per-unit fee to OEMs who agreed to minimum volumes exceeding expected shipments. The customer has no right under the contract to receive a credit for its unused units at the end of a contract. Nevertheless, Microsoft may allow the OEM to carry forward its unused licenses from the prior year Y.¹⁸ When Microsoft allows carry forward of unused copies, then the marginal price of MS-DOS in the current year is effectively reduced by its implicit rebate value.

In addition to the price incentives for exclusivity that are provided by the CPU license, Microsoft has been reported to have responded with a variety of direct penalties if an OEM shipped some of its machines with a competing operating system. First, the OEM may be prohibited from carrying forward unused MS-DOS licenses. At the extreme, Microsoft has on occasion required an OEM to renew the CPU license at equal or higher volumes to retain the carry-forward option. In this way, Microsoft's policy on carry forwards establishes a "tie" between sales from one year to the next.

¹⁶ See BERNSTEIN RESEARCH, *supra* note 4.

¹⁷ See *id.* at exhibit 3. In that year, Microsoft's sales of Windows through OEM and upgrades totaled \$599 million.

¹⁸ Whether the unit is marginal or inframarginal, its value is the reduction in *next year's* CPU license fee from displacing one unit—after discounting for time and likelihood that the additional unit will be used.

Second, Microsoft's technical service and support may be withheld from the rebellious OEM. This practice can disadvantage an OEM who needs this information to match the hardware configuration with the demands of the operating software (especially the choice of the microprocessor, the amount of RAM, and the graphics card).

Third, the price of Microsoft Windows has been increased to rebellious OEMs. As far back as the days when Microsoft Windows was called Interface Manager, Microsoft established a connection between the terms of sale of MS-DOS and its graphical interfaces.¹⁹ Microsoft cautioned OEMs against bundling competing multitasking interfaces (such as Quarterdeck's DESQview, VisiCorp's VisiOn and DRI's GEM) with PC hardware components such as hard disks.

Discounts on Microsoft Windows were extended to OEMs who agreed to accept a CPU license for MS-DOS. Those who refused the CPU license or who did not use MS-DOS exclusively, could still purchase Microsoft Windows. Again, Microsoft extended the nearly valueless option of purchasing it on an unbundled basis for a much higher per-unit price.

What options are open to an OEM who does not wish to exclusively ship its machines with MS-DOS? The OEM can negotiate a per-unit contract with Microsoft. However, Microsoft charges a price differential that is so high relative to CPU rates as to make the per-unit "option" economically infeasible. Alternatively, the OEM can choose not to deal at all with Microsoft. In that case, it can purchase OS/2 on a per-unit basis²⁰ (assuming the OEM manufactures machines that fit OS/2's higher memory requirements). Or it could send out machines with no operating system at all. A user who buys a "naked" machine must obtain an OS from another source. The owner can transfer the operating system from

¹⁹ See *supra* note 12.

²⁰ The CPU license appears to be unique to Microsoft. Besides DRI's per-unit license, UNIX is sold to OEMs such as Sun Microsystems using right-to-use "site" licenses. Like most site licenses, they provide for volume discounts.

an old machine, buy a new copy from a retail outlet, or "pirate" one from another user or an electronic bulletin board.

2. TECHNICAL INCOMPATIBILITIES Coordination on technical standards is crucial between the OS developer and applications developers. In several instances, Microsoft made it difficult for competitors, especially DRI/Novell's DR-DOS, to achieve compatibility with Microsoft Windows. Nowhere is this coordination more important than with the publication of the APIs. Microsoft has left undocumented some of these interfaces. In principle, access to these APIs would allow Microsoft to write applications (such as for its Word word processor or its Excel spreadsheet) that work faster and with greater functionality. Furthermore, should an applications developer discover and choose to use these undocumented interfaces, as long as they remain "unofficial," Microsoft can remove or alter them in later versions of the operating software, rendering parts of the applications useless.

One way for applications programmers to insure compatibility with an operating system is to receive copies of the preliminary version of the software. Known as "beta testing," this gives applications developers an opportunity to fine tune the interaction between the two programs.

In a well-publicized episode, DRI was excluded from the beta testing of Microsoft Windows 3.1 and later Microsoft's Windows for Workgroups product. The importance of compatibility testing with the Microsoft Windows beta version became evident when applications developers using DR-DOS received error messages warning them of a potential incompatibility with Microsoft Windows. As it turned out, upon installation, Microsoft Windows 3.1 checked whether the source of the underlying system and the extended memory manager was a Microsoft product. If they were not, the user was informed that a problem was detected, and was asked to contact Microsoft's beta support for Microsoft Windows 3.1. This message appeared on the screen even though no actual compatibility problem was detected. Indeed, if the user continued past the alleged error message, he or she would discover that Microsoft Windows 3.1 would run in conjunction with DR-DOS.

"The only error was that the customer was running Microsoft Windows on a competitor's version of DOS."²¹ The error messages raised fears of incompatibility among developers and users who contemplated running Microsoft Windows with non-Microsoft versions of DOS. Concerns over possible incompatibility between DR-DOS and Microsoft Windows resulted in significant declines in DR-DOS retail sales. In addition, Microsoft Windows disks included a "Readme" text file that cautioned users that "running Microsoft Windows 3.1 with an operating system other than MS-DOS could cause unexpected results or poor performance."²² Microsoft refused to address compatibility problems with DRI.²³ Microsoft boldly defended its action claiming it had no responsibility to assist an operating systems competitor.²⁴ Microsoft's actions went beyond refusal to assist a competitor, however, as it had engaged in commercial sabotage.

3. ANTITRUST ACTION Microsoft's practices first came to the attention of antitrust authorities in Korea. The Korean Fair Trading Commission launched an investigation that centered on use of the CPU license in Asia. In May 1992, the Korean FTC banned the use of CPU licenses in that country.²⁵ That action was not very effective, however, because Microsoft then began offering customer-specific price schedules with steep "cliffs" (sharp price reductions) at volumes close to the customer's requirements.

In June 1990, the U.S. Federal Trade Commission initiated a nonpublic (sic) investigation of Microsoft's practices. At first the focus of its inquiry was the relationship between Microsoft and

²¹ See *supra* note 12.

²² Goldman Rohm, Will the FTC Come to Its Senses About Microsoft's Mischief?, *UPSIDE*, August 1993, at 11-27.

²³ John Dodge, *It's Not an Uplifting Sight When Microsoft Bares its Claws*, *P.C. WEEK*, Dec. 9, 1991.

²⁴ Microsoft's decision was clearly related to market power. It provided Windows beta versions to many of its competitors in applications software where, at the time, it was not a dominant firm.

²⁵ Phang, *Microsoft Deals Not Fair: Korea*, *ASIA COMPUTER WEEKLY*, May 11-17, 1992.

IBM and the statements being made by these companies about their future OS product development. Later the investigation turned to marketing practices of DOS and Microsoft Windows.²⁶ Without ever acknowledging the investigation, the Commission met in February 1993 to decide whether to issue an unprecedented preliminary injunction requiring Microsoft to cease and desist from its marketing practices. The vote was a 2-2 tie. Six months later, the FTC deadlocked again, apparently dooming any chance of antitrust action by the U.S. government.

But then, with victory almost within Microsoft's grasp, in an unprecedented move, the Antitrust Division of the Department of Justice (the Department) took up the case and, after extensive further investigation,²⁷ negotiated a consent decree with Microsoft. On July 15, 1994 the Department filed a civil antitrust complaint along with a proposed Final Judgment to which Microsoft had consented (the Consent Decree),²⁸ followed, as required under the Tunney Act, by a Competitive Impact Statement (CIS).²⁹ The case then took an even more startling twist when Judge Sporkin of the DC District Court refused to play the role of a mushroom³⁰ and rejected the decree as inadequate under the Tunney Act. The U.S. government and Microsoft jointly (again, an unusual, event) appealed Judge Sporkin's decision, and that appeal is currently

²⁶ At one time, the FTC staff was also investigating whether the relationship between Microsoft's operating systems and applications divisions created remediable competitive problems in markets for applications software.

²⁷ AAG Bingaman revealed that the investigation had consumed thousands of hours of attorney and economist time.

²⁸ *United States v. Microsoft Corp.*, No. 94-1564 (D.D.C. filed July 15, 1995). Amended versions of the Proposed Final Judgment and the Competitive Impact Statement were filed with the court on July 27, 1994.

²⁹ Proposed Final Judgment and Competitive Impact Statement, 59 Fed. Reg. 42845 (1994) (proposed August 19, 1994).

³⁰ The court noted that "Tunney Act courts are not mushrooms to be placed in a dark corner and sprinkled with fertilizer." *Microsoft*, 1995 U.S. Dist. LEXIS 1654, at 42.

pending before the U.S. Court of Appeals for the District of Columbia.

The proposed consent decree is described and evaluated below in section VI. Before doing so, however, we first examine two potential efficiency rationales for the CPU licenses—the apparent similarity between the CPU license and an efficient two-part tariff, discussed in section III, and the argument that the CPU license may be an effective way to combat fraud and piracy, discussed in section IV. We discuss the potential anticompetitive rationales for Microsoft's practices in section V.

III. First-degree price discrimination vs. inefficient substitution

At first glance it may appear that the CPU license is just a means to provide volume discounts to large OEMs. This is not the case: it is possible that an OEM that purchases *more* MS-DOS pays a *higher* per-unit price than one that purchases *fewer* units. This would happen if an OEM purchased more units of MS-DOS than some another OEM, but proceeded to ship many more machines (loaded with an alternative DOS or none at all). It is possible that its per-unit cost of MS-DOS could be higher because its royalty charge is based on the number of machines shipped.³¹

In fact, CPU licenses may imply a quantity *premium* if units beyond the requirements are sold at a per-unit charge. At the minimum requirements quantity, the marginal price jumps from zero to a positive level. So based on marginal prices, purchases beyond the requirements level incur a quantity premium.³²

³¹ The actual price paid per unit could be higher even if the royalty fee itself incorporated volume discounts.

³² Average price is the more typical yardstick for measuring nonlinearity of prices. In the case of a CPU license, they fall through the range up to the minimum requirements and thereafter may rise or fall depending on whether the average price at the requirements level is lower or higher, respectively, than the per-unit charge for additional sales.

The CPU license could be characterized as first-degree price discrimination only in a very limited sense. An OEM who signs under the CPU license (or a take-or-pay license with $X > \text{output}$) has agreed to a lump-sum payment, with an (expected) zero marginal price for 1 year. However, since the size of the lump-sum payment is based on expected sales times a per-unit royalty, the OEM knows that if its sales increase, the (apparent) lump-sum payment next year will also increase proportionately (based on the per-unit royalty Microsoft will be charging in the next year). Thus, for any time horizon longer than 1 year, the CPU license is a tax on output; it is not first-degree price discrimination.

A. First-degree price discrimination

But even if—perhaps within the very limited time horizon of a 1-year window—OEMs regarded the CPU license as imposing a lump-sum fee unrelated to their MS-DOS use or to their output, such a royalty structure would be neither profitable to Microsoft nor would it be welfare-enhancing when compared to a per-unit royalty.

Efficient first-degree price discrimination occurs when a seller charges a two-part fee, consisting of a lump-sum payment for the right to purchase the product and a price for each unit equal to marginal cost. Where, as here, marginal cost is essentially zero, first-degree price discrimination requires a per-unit price of zero.

Economists have long recognized the strong efficiency advantages of first-degree price discrimination when customers are final consumers so that their demands are independent. But these results do not hold when intermediate inputs are sold to competing downstream firms. In that case the demands of such customers (the firms in the downstream industry) are clearly not independent (i.e., if my rivals pay less for an input than I do, the price of the final product falls, reducing my demand for the input). Ordover and Panzar state the issue quite clearly.

... we recast the welfare analysis of the simple two-part tariff using the classical model of perfect competition in which all firms are identical and free entry and exit ensures that the equilibrium output price

is equal to minimum average cost. In this context we discover that two-part tariffs are not generally desirable from a welfare standpoint. . . . This is due to the fact that the entry fee, instead of acting as a "lump sum levy," affects both the equilibrium number of firms and their output level. This new distortion must be balanced against the losses due to a unit price in excess of marginal cost.³³

However, where, as here, the input (OS) is used in fixed proportions with the output (PCs), and the downstream industry is a classic competitive industry with U-shaped average cost curves, Ordover and Panzar find that a very strong theoretical result obtains: a monopoly seller of the input would find any two-part tariff, including an all-or-nothing arrangement where marginal cost to the buyer is zero, *less profitable than a uniform per-unit fee*. In addition, the uniform per-unit fee results in *higher economic welfare than any two-part tariff*. As Ordover and Panzar put it,

Most surprisingly, for the empirically relevant class of production processes in which the purchased input is required in fixed proportion to output, we discover that a two-part tariff is never optimal from either a profit or welfare maximizing standpoint (at 660).

The intuition behind this result is rather straightforward. It is well known that under fixed proportions an upstream uniform pricing monopolist can extract all the profits which an integrated uniform pricing monopolist could reap. Since competition downstream ensures that a uniform price prevails in the final product market, there can be nothing to gain from introducing a two-part tariff; optimal choice of [the per-unit price] allows the monopolist to earn the maximum possible under such circumstances. There is something to lose, however, since an entry fee $e > 0$ causes the downstream firms to operate at an inefficiently large scale. Total (upstream plus downstream) costs are not minimized and a portion of this dead-weight burden falls on the monopolist. Viewed another way, this result reveals the futility of attempting to impose a seemingly nondistortionary lump-sum levy e on a perfectly competitive industry with free entry and exit (at 666-67).

³³ See Ordover & Panzar, *On the Nonlinear Pricing of Inputs*, INT'L ECON. REV., October 1982, at 659-60.

In short, even if Microsoft's CPU license (or equivalent volume discounts) did impose a true lump-sum payment, there would be no efficiency or welfare gain that could provide a defense for such a pricing system. Nor would such a licensing system be profitable for Microsoft to impose on OEMs even if those OEMs somehow did not recognize the link between their sales and the lump-sum royalties they paid. Both theory and the available evidence would indicate, therefore, that Microsoft's CPU license (or its equivalent in the form of a volume discount) is not a form of first-degree price discrimination.³⁴

One might ask if, in the context of this case, the fixed-proportions assumption made by Ordover and Panzar refers to fixed proportions between the OS and the PC, or does the result require fixed proportions between MS-DOS and the PC?³⁵ It does seem safe to assume that every PC requires one (and, as a practical

³⁴ It could be argued that the Ordover and Panzar analysis is too long run, and that what Microsoft is really attempting to do is expropriate some of the fixed, sunk costs OEMs have committed. There is no obvious gain in economic efficiency from such behavior. More importantly, for Microsoft to do this, it would have to (1) estimate the fixed sunk costs of every manufacturer at the time DOS was first offered, and set a different lump-sum fee for every manufacturer that was less than that manufacturer's sunk costs; (2) convince the manufacturer that the lump-sum fee would not change if that manufacturer produced fewer or more PCs; (3) set a zero license fee for DOS to any manufacturer who entered the PC market after DOS came on the OS market; and (4) set an average price for MS-DOS to the more established, larger OEMs that have expended significant sunk costs in differentiating their product that is significantly higher than the average price to smaller "generic" OEMs, such as, notably, many of the Taiwanese and Korean OEMs. Conditions (1) and (2) would appear to be difficult to achieve. Condition (3) does not hold since Microsoft has not waived the CPU requirement for new OEM entrants. Finally, Microsoft appears to be charging the larger, more successful OEMs a lower—rather than higher—average price for MS-DOS in contradiction with condition (4). Thus, this possible argument for Microsoft appears to have no merit, either on policy or factual grounds.

³⁵ Note that fixed proportions implies that the input is essential, but an input can be essential without fixed proportions: an input is not essential if the unit cost of output approaches a finite limit as the price of the input approaches infinity.

matter, only one) OS. Every PC may not, however, require one MS-DOS although again, as a practical matter, given the pricing and tying arrangements imposed by Microsoft, most OEMs appear to have no practical alternative to MS-DOS over any relevant price range.

Recall, however, how the fixed-proportions assumption affects the results of the model. Absent fixed-proportions, the lump-sum fee is still inefficient, it is just that the per-unit-MS-DOS royalty alternative becomes less desirable. The analysis of the effect of the lump-sum royalty is only affected by alternatives to MS-DOS if those alternatives were such close substitutes that an increase in the lump-sum fee would not impose any cost on the OEM. That is, a lump-sum fee of any size would cause all OEMs to switch over to the alternative OS.

With fixed-proportions of MS-DOS per PC, a per-unit-MS-DOS royalty would not affect the MS-DOS to PC ratio. Thus a positive fee per unit of MS-DOS would not result in inefficient input proportions, making a per-unit-MS-DOS royalty as socially efficient (or inefficient) as a per-unit-PC royalty. This is in contrast to a lump-sum fee that would lead to, in the final market equilibrium, a distorted ratio of fixed to variable inputs, with too few firms, each producing too much output.

Suppose, instead, that there are fixed proportions between having some OS and a PC, but not having MS-DOS and a PC. In that case, a per-unit-MS-DOS royalty could result in the OEM shifting to some other OS. Would this be inefficient? Only if the social cost of the OS alternative (for example, DR-DOS) were higher than the social marginal cost of MS-DOS. But since the social marginal cost of both DR-DOS and MS-DOS (and indeed of all OSs) are equal (and nearly zero), this change in the ratio of MS-DOS per PC does not introduce an inefficiency.

We thus conclude that sufficient conditions for the Ordover and Panzar result are that either (a) the particular input—in this case MS-DOS—is used in fixed proportions with the final product, or (b) the social marginal cost of the alternative DOS chosen

in response to an increase in the per-unit price of MS-DOS is less than the social marginal cost of MS-DOS.

This implies a natural extension of the Panzar-Ordover result. Suppose that the quality-adjusted social marginal cost of the input(s) B, C, . . . toward which a manufacturer would substitute if the price of input A were increased is less than the quality-adjusted social marginal cost of input A. This appears to be true in the case of MS-DOS and DR-DOS since the production cost of both MS-DOS and DR-DOS are the same, while DR-DOS is of arguably higher quality than MS-DOS. Production efficiency will then actually increase as a result of the higher per-unit-price for MS-DOS. It follows that, under these circumstances, a per-unit-MS-DOS royalty would be socially even more efficient than a per-unit-PC royalty.

B. Inefficient substitution

While the CPU license does not produce a positive output effect (i.e., encourage efficient utilization of a zero-marginal-cost input), it does have a significant substitution effect. The CPU license induces substitution of MS-DOS for OS alternatives. While this may be privately profitable, the social gain is zero, even if it did not induce the exit of rival operating systems such as DR-DOS with its attendant expected effects on raising the MS-DOS license fees. Both MS-DOS and any other OS have a near zero social marginal cost in use. Thus, to the extent that the CPU license induces substitution of MS-DOS for DR-DOS, no cost saving results.³⁶ Indeed, if, as appears to be the case, the value of DR-DOS is greater than that of MS-DOS (at least on those PCs where DR-DOS would be incorporated, absent the CPU license for MS-DOS), the substitution of MS-DOS for DR-DOS actually reduces efficiency and total welfare.

³⁶ This is, of course, just another example of the theory of the second-best.

IV. Antifraud and antipiracy rationales for imposing CPU royalties rather than MS-DOS unit royalties

In this section we compare CPU royalties and per-unit royalties in terms of their effects on fraud (underreporting of volume by OEMs) and piracy (unauthorized copying by end-users).

A. Reducing manufacturer fraud

As a factual matter, we are not aware of any evidence that underreporting of MS-DOS usage by PC manufacturers is or has ever been a serious problem. This is not to say however that in principle a CPU license could not reduce fraud. OS systems are duplicated and incorporated by the manufacturer, not the licensor, and only the manufacturer knows how many units have been duplicated. If the licensor could readily determine the total PC production of the manufacturer, a CPU license could effectively prevent such fraud. Several considerations, however, point to manufacturer fraud as not being the rationale for CPU licensing by Microsoft.

1. First and most important, we should observe CPU licenses where fraud is likely, but where market power or exclusion could not plausibly provide a rationale. If CPU licenses are not observed in those circumstances, then the ability of CPU licenses to minimize fraud is highly suspect.

Microsoft has historically chosen not to license many smaller and/or "high-risk" OEMs, i.e., OEMs (mostly in Asia-Pacific) that are viewed as particularly prone to engage in fraud or whose demand levels might not justify the fixed costs of licensing. By contrast, DRI did license many of these producers, employing a variety of means to control potential fraudulent underreporting. These methods include the use of serial numbers cross-referenced to end-user requests for technical support, the use of holograms, and audits of the OEMs. If CPU licenses were an effective and inexpensive means of controlling fraud, DRI might well have used CPU licenses for these OEMs.³⁷ Beyond this, fraud is not a

³⁷ DRI did not believe that the expense of an audit was affected, one way or the other, by whether one is auditing for violation of a CPU

problem limited to operating system vendors, yet to our knowledge other software vendors similarly situated have not used such licensing terms. Thus, this market test indicates that CPU licenses are not part of the optimal antifraud strategy for a licensor of operating systems where market power and exclusionary potential are not present.

2. Fraudulent underreporting cannot credibly be asserted for every OEM. Especially at the currently low royalty rates that would continue with competition from DR-DOS, few, if any, of the major OEMs would have any incentive to defraud MS. Even if detection by Microsoft were unlikely, an OEM can be deterred³⁸ from fraud by penalties that are a correspondingly large multiple of the gains from fraud; Microsoft could clearly impose a penalty on a major OEM that would dwarf that OEM's potential gain from fraud. This implies that Microsoft imposes CPU licenses on many OEMs that could not credibly be expected to attempt to defraud Microsoft.

3. To our knowledge, Microsoft does not regularly "audit" the number of PCs produced. Rather, Microsoft relies on that manufacturer's representations of the number of PCs, just as they might have to rely on that manufacturer's representations of the number of PCs produced that incorporated MS-DOS.

One might argue that if an individual OEM's shipments of PCs can be accurately estimated, the use of CPU licenses reduces the need to engage in audits even if, given the decision to audit, there is no difference in auditing costs between CPU licenses and per-unit licenses. Given that public information on OEM shipments shows considerable variation in estimated shipments for many OEMs, however, there is reason to question the empirical significance of such an argument.

There is, however, a plausible way to test this assertion, using information that, while not privately available, would be available

license or a per OS-unit license. Its business behavior confirms this belief.

³⁸ For example, the contract could specify treble damages (a three-fold royalty) for all proven cases of underreporting.

to an enforcement agency such as the FTC or DOJ. What has happened to Microsoft's auditing expenses (per MS-DOS unit) since it introduced CPU licenses? If auditing expenses did not fall, this is inconsistent with CPU licenses being an inherently superior method for dealing with OEM fraud.³⁹

4. The timing of any shift between MS-DOS unit licenses and CPU licenses should parallel shifts in the seriousness of OEM fraud. To our knowledge, however, the introduction of CPU licenses did not coincide with any notable increase in the incidence or the potential for OEM fraud.

In particular, since an OEM's incentive to underreport is proportional to the size of the royalty, the extent of CPU licensing should have risen or fallen with the Microsoft royalty. This implies that, since the Microsoft royalty rate rose steadily until the availability of DR-DOS induced a sharp decline, the prevalence of CPU licensing should have increased over time until DR-DOS appeared, and then declined. In contrast, it is our understanding that widespread CPU licensing did not appear until early 1990, after the introduction of DR-DOS and the ensuing decline in prices for PC operating systems.

5. OEM fraud could be handled by a credited CPU license, with the royalty based on the number of PCs minus the number of documented machines that were sent out either naked or loaded with some alternative OS system. Yet Microsoft's CPU licenses do not work in this fashion. Instead they require the OEM to pay a royalty on all machines produced, without a credit for the use of an alternative DOS.

6. Finally since, as noted above, an OEM's incentive to underreport is proportional to the size of the royalty, and since the availability of DR-DOS has resulted in a sharp decline in Microsoft's royalty, it would appear that maintaining effective competi-

³⁹ The test is only a one-way test. If auditing expenses fell, this could be due to the reduction in the incentive for fraud (and, therefore, the reduced need to spend money to control it) after competition drove down the price of MS-DOS. This is discussed in points 5 and 6 below.

tion in the OS market has a desirable by-product—it reduces the extent of OEM fraud.

B. Reducing number of naked machines in order to reduce piracy by the OEM's customers

Piracy of software by computer users has long been a concern of software developers. It has been alleged that CPU licenses may discourage piracy by reducing the relative benefits of installing an unlicensed copy of MS-DOS. By not installing any operating system on an outgoing machine, an OEM could pass along the saving to its customers. The OEM would sell more machines because at least some customers will prefer to install a pirated OS (and comparatively higher cost) rather than pay the additional cost to the OEM. This option is no longer attractive when the OEM faces a zero price for installing MS-DOS on a machine as is true under a CPU license.

There are, however, several difficulties with this proposition.⁴⁰

1. Even before the advent of the CPU license, virtually no computers from companies licensed by Microsoft were shipped naked (i.e., without an OS).⁴¹ Simply put, if the incidence of naked machines shipped by OEMs licensed by Microsoft was very low before CPU licenses, then it is not plausible that CPU licensing was implemented to reduce the shipping of naked machines.

2. Another major factor is that the OEM license fee is far lower than the retail price of MS-DOS. It might seem reasonable that individual users would pirate copies themselves (or buy commercially pirated copies of MS-DOS) if their alternative were a

⁴⁰ Note that the only form of piracy relevant here is an end-user's substitution of pirated MS-DOS for the MS-DOS incorporated by an OEM. In what follows, therefore, "piracy" refers to only this substitution, and not to the copying of upgraded versions of MS-DOS by customers whose machines incorporated earlier version of MS-DOS. Such "retail upgrade" thievery could not be affected by a CPU license.

⁴¹ This is not surprising, since there are good reasons to believe that it would be more expensive for most end-users to install an OS themselves.

retail MS-DOS at a price of \$79. But it seems inherently unlikely that many buyers of a new computer, no matter how much they may dislike paying for software, would go to such lengths to avoid paying less than \$15 more on a \$1000 computer system.

3. A CPU license could only deter piracy if, absent the CPU license, the OEM would provide a discount for machines without an OS. If the customer pays the same price for a naked machine as for one with an OS installed, purchasing a naked machine cannot facilitate piracy. It is our understanding, however, that OEMs did not generally offer a discount for naked machines prior to the introduction of the CPU license. If correct, then the CPU license could not have been a response to piracy. Moreover, even if prior to the CPU license, Microsoft's OEMs had been offering naked machines at a discount to end-users, Microsoft could easily have handled the problem by forbidding its OEMs from offering a discount for naked machines. While this might conceivably have raised some antitrust risk for Microsoft (under some private plaintiff's theory that the vertical restraint was unlawful), that antitrust risk would surely be far less than that associated with the CPU license. Thus, fear of antitrust exposure could not explain Microsoft's failure to adopt this simpler remedy, which would eliminate any incentive for end-user piracy while not foreclosing OS competition.

4. Microsoft could prevent the problem, if it exists, by charging a royalty based on the number of PCs that were sent out either with MS-DOS or without any OS or, if OEM fraud is also a problem, imposing a royalty based on the number of total PCs minus the documented number of machines that were sent out incorporating some alternative OS.

5. Since the customer's incentive to pirate will be proportionate to the amount of the OEM royalty that is passed on to him, it also follows that, as in the case of OEM fraud, a customer's incentive to pirate an OS will be proportional to the size of the OEM royalty. It would thus appear, again, that maintaining effective competition in the OEM market for operating systems has the desirable by-product of reducing the incentive for piracy.

6. If Microsoft is earning supracompetitive returns (i.e., risk-adjusted returns above what would be necessary for the product to have been developed), piracy may be socially beneficial if it results in greater use of MS-DOS than would otherwise be the case. Output increases to some extent because some users of pirated MS-DOS might not otherwise have bought MS-DOS. A more important effect, however, is on the demand for MS-DOS becoming more elastic if piracy is possible since piracy is an increasing function of the price charged the OEM for MS-DOS. Thus even "deserving" customers—those who would not or could not pirate an OS system—gain from the existence of the pirates.⁴²

Microsoft may object that tolerating piracy would inefficiently interfere with its ability to recover its investment in developing intellectual property. However, it is simply not the case that efficiency is greatest if property rights in intellectual property, at least as those rights are currently established, are perfectly secure. Most inventors are free-riding on a common pool of prior knowledge. In such circumstances, it can be shown that it is optimal for the inventor to face the same degree of free-riding on her invention. Perfect appropriability, in such circumstances, would lead to overinvestment in inventive activity.⁴³ This observation seems particularly relevant here. Microsoft's original MS-DOS was

⁴² The situation is reversed, of course, in a competitive market, where pirating would only increase costs to nonpirating customers or even make some products unavailable.

⁴³ See Barzel's qualification (Yoram Barzel, *Optimum Timing of Inventions*, 50 REV. ECON. & STAT. 348 (Aug. 1968)) of Arrow's argument (Kenneth I. Arrow, *Economic Welfare and the Allocation of Resources for Invention*, NATIONAL BUREAU OF ECONOMIC RESEARCH, THE RATE AND DIRECTION OF INVENTIVE ACTIVITY (1962) that imperfect appropriability leads to underinvestment in inventive activity. For the proposition that the problems raised by imperfect appropriability may be significantly reduced if the existence of the product results in pecuniary and/or technological effects, see Jack Hirshleifer, *The Private and Social Value of Information and the Return to Inventive Activity*, 61 AM. ECON. REV. 561 (Sept. 1971). See also A. Michael Spence, *The Economics of Internal Organization: An introduction*, 6 BELL J. ECON. 163, 168 (Spring 1975), for a critique of Arrow.

based on its acquisition of technology from a company that had cloned CP/M, a pioneering PC operating system developed by DRI.

Thus we have an interesting policy paradox. The claim that CPU licenses help to control piracy is highly questionable when Microsoft earns monopoly returns on MS-DOS and there appears to be no efficiency rationale for the supranormal returns. Were Microsoft earning competitive returns to MS-DOS, controlling piracy would clearly be defensible on efficiency grounds. Under the CPU license, however, the act that (allegedly) controls piracy also forecloses competition, thus also undermining the policy rationale for controlling the (alleged) piracy.

7. A similar debate over copying in other industries has made it clear that the effects of copying on profits and on social welfare are complex and ambiguous. The existence of "unauthorized" copying may actually increase the profits to the seller, and may be socially desirable, depending on the nature of infringing users and the heterogeneity of purchasers.⁴⁴ It is simply not clear that the

⁴⁴ To illustrate the point, assume 1000 homogeneous buyers, each with a reservation price for DOS of \$100 for her main machine. Each buyer also has a second PC, onto which she can copy her purchased copy of DOS, for which the reservation price for DOS is \$40 (alternatively, she has a younger brother, partner in her office, etc., to whom a pirated version of DOS would be worth \$40). The marginal cost of DOS to Microsoft is \$10, and the marginal cost to a "pirate" of unauthorized copying DOS is \$15. If Microsoft effectively prevented copying (perhaps by inserting a secret virus that would destroy the disk if copied) it would set a price of \$100, selling 1000 units, with revenue of \$100,000 and profits of \$90,000. (If preventing copying did not cost less, profits would be less than \$90,000.) If Microsoft made no effort to prevent copying/piracy, it would set a price of \$125 ($= \$100 + \$40 - \15), selling 1000 units for a revenue of \$125,000 and profits of \$115,000. Note that, in this example, anything that increases the cost to the "pirate" of unauthorized copying reduces the profits to the licensee. The point of this example is not that all copying/piracy benefits the licensee, but rather that even the private, much less the social, effects of copyright violation are complex and often ambiguous. For a thorough analysis of these issues, see TIMOTHY J. BRENNAN, *TAXING HOME AUDIO TAPING* (Economics Analysis Group Discussion Paper EAG 86-5, Antitrust Division, U.S. Department of Justice (April 15, 1986) and the articles cited therein.